Distributed Protocols and Heterogeneous Trust: Technical Report

Isaac C. Sheff
Cornell University
isheff@cs.cornell.edu

Robbert van Renesse
Cornell University
rvr@cs.cornell.edu

Andrew C. Myers
Cornell University
andru@cs.cornell.edu

arXiv:1412.3136v1 [cs.DC] 9 Dec 2014

Abstract

The robustness of distributed systems is usually phrased in terms of the number of failures of certain types that they can withstand. However, these failure models are too crude to describe the different kinds of trust and expectations of participants in the modern world of complex, integrated systems extending across different owners, networks, and administrative domains. Modern systems often exist in an environment of heterogeneous trust, in which different participants may have different opinions about the trustworthiness of other nodes, and a single participant may consider other nodes to differ in their trustworthiness. We explore how to construct distributed protocols that meet the requirements of all participants, even in heterogeneous trust environments. The key to our approach is using lattice-based information flow to analyze and prove protocol properties. To demonstrate this approach, we show how two earlier distributed algorithms can be generalized to work in the presence of heterogeneous trust: first, Heterogeneous Fast Consensus, an adaptation of the earlier Bosco Fast Consensus protocol; and second, Nysiad, an algorithm for converting crash-tolerant protocols to be Byzantine-tolerant. Through simulations, we show that customizing a protocol to a heterogeneous trust configuration yields performance improvements over the conventional protocol designed for homogeneous trust.

1. Introduction

Fault tolerance is critical for distributed systems. Traditionally, distributed systems and protocols are designed around the ability to tolerate some number of failures, sometimes differentiated by type, such as crash or Byzantine [8, 20, 29, 30]. For well-studied problems such as consensus, lower bounds are traditionally expressed in terms of the number of participants needed to tolerate some number of failures f [6, 7, 20, 32]. In our increasingly interconnected world, however, systems must routinely operate across locations and between different owners, requiring a richer notion of what failures are possible. In complex systems integrated across administrative domains (that is, federated systems), different participants in the system may not even agree on what types of failures may occur or where in the system [4, 11, 16, 33].

For example, suppose Alice is building a new app, a competitor to Bob's. However, due to the nature of the application, it is best if Alice and Bob's apps agree (quickly) on some things. Alice's app works with servers maintained by Carol, Dave, and Eve. Due to strong records and reliably enforced SLAs, Alice is willing to allow her program to fail if Carol, Dave, or Eve do. She does not tolerate their failure. She fears, however, that Bob may lie to her. On the other hand, Bob, Carol, Dave, and Eve do not know Alice (her app is new, after all), and believe she may lie or crash. As established businesses with contracts and

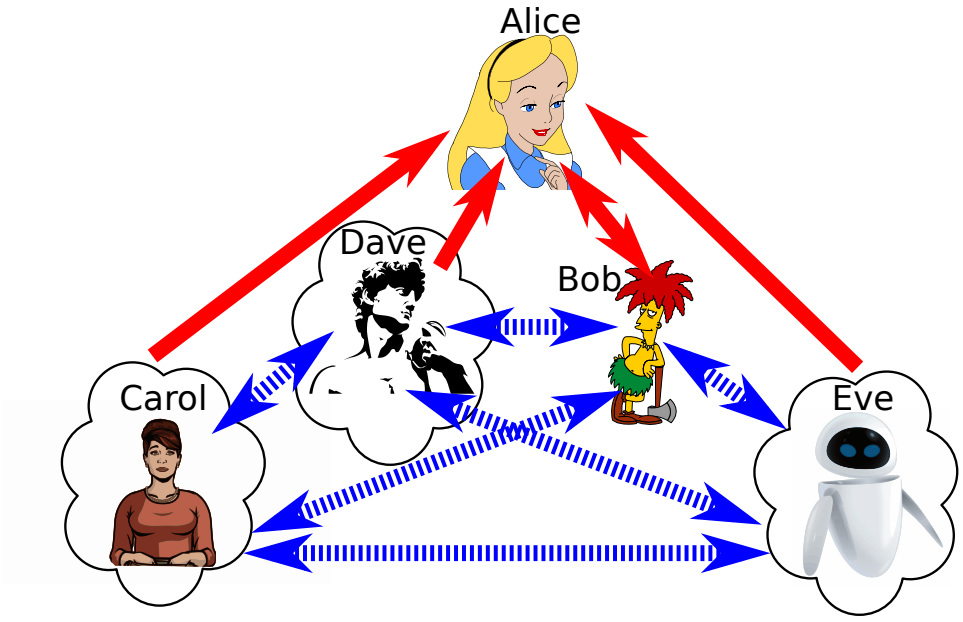

Figure 1. Red solid arrows represent participants' belief that another participant might lie (integrity failure), while blue striped arrows represent participants' belief that another participant may crash (availability failure). Everyone but Alice believes at most one crash may occur, and Alice may lie in addition. Alice believes only Bob can fail, although he may fail in any manner.

track records, however, they believe each other to be honest, although they have a healthy tolerance for at most one crash among themselves.

This trust configuration, depicted in Figure 1, is much more complex than in the traditional models where participants uniformly agree on the maximum number of crash failures and the maximum number of Byzantine failures. In this paper, we explore the possibility of more general distributed protocols that take into account and even exploit such complex, heterogeneous trust.

1.1 Contribution

Heterogeneous trust presents a challenge for designing fault-tolerant protocols. Our key idea is to use information flow methods to reason about the integrity and availability properties of distributed systems. Prior work on information flow methods has mostly addressed confidentiality properties of systems [28], though some prior work [35, 36, 38, 40] has addressed integrity and availability properties in a limited way. However, our work exploits information flow in a more general and sophisticated way.

We construct and analyze fault tolerant protocols by tracking the integrity and availability of flows of information through the protocol. Intuitively, each participant characterizes its assumptions about the availability and integrity of system components as a *label* drawn from a rich lattice that expresses the trust of participants. The labels are expressive enough to represent various combinations of crash and/or Byzantine failures, as expected by the various participants. Using these labels, it is then possible to analyze under what conditions the results of running a protocol have an availability and integrity acceptable to each participant.

When participants have different opinions about the trustworthiness of the system components, a new phenomenon arises: some participants' trust assumptions may be violated while others' assumptions still hold. We consider a protocol to be correct only if the violation of one participant's trust assumptions cannot damage the availability or integrity of the system as viewed by any participant whose assumptions have not been violated. Naturally, not all protocols can be run with all possible configurations of trust among participants; each protocol has minimal requirements concerning those participants' trust in each other.

As an example of this style of synthesis, we present Heterogeneous Fast Consensus, a generalization of the "Bosco" protocol [32]. This generalization achieves the same bounds when traditional homogeneous trust assumptions are used, but the protocol is capable of operating with heterogeneous assumptions. We generalize the traditional properties of a consensus protocol—Unanimity, Validity, Progress, and Termination—for heterogeneous trust environments [6, 32]. We explain this protocol's requirements, and prove it satisfies these properties. Using simulations, we also demonstrate that Heterogeneous Fast Consensus offers significant advantages in speed and resource requirements when compared to the Fast Consensus protocol it generalizes.

As a second example, we develop a generalization of Byzantine tolerant Ordered, Asynchronous, Reliable Broadcast (OARcast) and Nysiad, an algorithm for converting crash-tolerant protocols into Byzantine-tolerant protocols [13, 14]. This generalization demonstrates that in some cases, Availability and Integrity can affect each other in counterintuitive ways. For example, Nysiad includes cases in which a participant cannot affect a value's integrity, but can make it unavailable by lying.

1.2 Related Work

Others have looked into richer notions of failure, including generalizing f-failures to failure-prone and survivor sets [16], or mixing availability (*omission*, or *Crash*) and integrity (*commission*, or *Byzantine*) failures [12, 23, 31]. Some work operates on expanded failure models, including notions of selfish but not malicious participants, sometimes mixed with other kinds of failure [1, 3]. A distinctive feature of our work is the removal of the assumption that all participants share the same notion of possible failures. Properties of the form "as long as the failure assumption is not violated, it is guaranteed that ..." can be generalized to explain which guarantees can be made for which participants.

This work exists at the intersection of information flow analysis, a technique traditionally applied through programming languages (e.g., [25]), and distributed systems theory. Most prior research on language-based security has not concerned itself with fault tolerance, but there are exceptions. Zheng has explored using information flow to reason about availability [38] and integrity [40] properties, including in distributed systems [37, 39], but the focus has been on linguistic mechanisms and simple quorum-based protocols. Walker et al. [34] design a lambda calculus formalizing the possibility of an integrity fault on a single machine.

Consensus, of course, is a widely-studied topic under a variety of failure models [15, 18, 19, 32]. Our particular generalization of the Bosco Byzantine fast consensus protocol [32] serves as the first known example of consensus under heterogeneous mutual distrust.

2. System model

2.1 Network

Our trust model is similar that used in other models of distributed systems [18, 30, 32]. We assume an asynchronous network environment in which any participant can communicate with any other. There is no guarantee, however, that participants trust each other. We assume that the network is, or can be made, reliable. That is to say, our participants are assumed to have whatever message resending protocols (e.g., [2, 27]) are necessary to guarantee that the only case in which a message never arrives is the case in which the sender or receiver have failed in some way. Additionally, we assume that faulty participants cannot forge the source identity of a message sent by a correct participant; cryptographic signatures are one way to achieve this.

In order to guarantee probabilistic termination in Heterogeneous Fast Consensus, we also assume that for any set of messages, no one of which is causally before another, if the set is sent repeatedly, the network will eventually deliver them in each possible order.

2.2 Heterogeneous Trust

Each participant has its own assumptions concerning the availability of the system (who might crash), and the integrity of the system (who might lie, or do something other than correctly execute the protocol). These assumptions can be thought of as describing what "attacker" is expected by each protocol participant. The limits of this attacker's power can be captured using information-flow labels (See section 3). In a system with heterogneous trust, participants that have not failed can then be characterized as either **gurus** or **chumps**:

Gurus are participants who function correctly and whose trust assumptions are not violated. By definition, no set of failures that actually occur can violate the availability or integrity expectations of a guru. For most protocols, most guarantees made pertain to gurus.

Chumps are participants who function correctly—they obey the prescribed protocol and do not crash—but whose trust assumptions have been violated. In traditional failure-tolerant systems with homogeneous trust, all participants make the same trust assumptions, so either all correct participants are chumps (in which case few or no guarantees are made), or everyone is a guru. Here, we must be more nuanced. Unsurprisingly, chumps may receive "wrong" results.

Recall the example from section 1. Five participants have a trust configuration in which everyone but Alice tolerates one crash among themselves, and tolerates Byzantine behavior from Alice. Alice tolerates failures only from Bob, but tolerates Byzantine failures on his part.

Suppose the participants wish to achieve consensus on data for Alice's app quickly, ideally in a single round in the usual case. For example, they might want to use the Bosco Byzantine consensus protocol, which can achieve consensus quickly. For Bosco, however, at least nine participants would be required, since some participants believe at least one Byzantine failure can occur, in addition to one crash. The five current participants would have to recruit others with relevant trust properties. In fact, tolerating just one Byzantine failure requires at least 6 participants for any one-communication-step consensus protocol tolerating f

failures [32]. As we show in section 5, it is nevertheless possible to create a variant of the Bosco fast consensus protocol that satisfies the requirements of all five participants, with *no* additional participants.

3. Information flow policies

3.1 Labels

Information flow control offers a way to reason about the properties of information in a system. While most prior work on information flow has concentrated on proving confidentiality properties, it is also possible to reason about the integrity [24, 26] and availability [38] of information. In this work, we focus on availability and integrity of the information used in distributed protocols, and leave confidentiality concerns to future work.

To support the analysis of information flow, all information in the system is assigned a label drawn from a lattice of labels that express information security requirements for the labeled information. [10, 24]. As information flows through the system, its label (ordinarily) moves only upward in the lattice. With dynamic information flow control, this label is represented explicitly at run time, whereas with static information flow control, it is merely a compile-time aspect of the information. In this work, we use static information flow control, because we want to design protocols whose properties are verified before execution.

In this work, we adapt the Decentralized Label Model (DLM) [24] to capture the integrity and availability requirements of information used in protocols. The DLM is designed for systems in which principals are mutually distrusting, which is ideal for the design of distributed protocols.

3.2 Principals

Policies are expressed in terms of *principals*, which may represent machines, users or other entities to whom permissions may be given. One principal may be trusted at least as much as another principal. If a principal p is at least as trusted as a principal q, we say that p *acts for* q, written $p \succcurlyeq q$.[1] Any action that can be taken by q can also be taken by p, meaning that p can act with the full authority of q. The universally trusted principal $\top$ can act for everyone; $\bot$ is the principal with minimal authority, for whom everyone may act.

Compound principals are a way of expressing principals representing the actions of multiple principals [21, 38]. In particular, we use the *conjunctive principal* $p \wedge q$ to represent the least upper bound of the authority of p and q (essentially, their combined authority) and the *disjunctive principal* $p \vee q$ to represent their greatest lower bound. Therefore, for any principals p and q, we have $p \wedge q \succeq p \succeq p \vee q$. Formal rules for compound $\succeq$ can be found in Appendix A.

3.3 Policies

A label is a set of policies, and the label is enforced exactly when all the policies in ℓ are enforced. A policy is a statement that some principal, an owner, trusts some other principal to affect the labeled information. Two kinds of policies are considered here. A policy of the form $o \xleftarrow{I} p$ signifies that owner principal o trusts only principal p (or other principals that act for p) to affect the content of this

[1] This is essentially the same idea as "speaks for" [21] in authorization logics.

information; similarly policy $o \overset{A}{\leftarrow} p$ means that o trusts p with the availability of this information (that is, trusts p to not make it unavailable).

This makes intuitive sense if p is a host principal, but p could also be a compound principal representing multiple hosts. For example, a policy $o \overset{I}{\leftarrow} p \vee q$ means that o trusts both hosts p and q with the integrity of the labeled data; a failure of either principal can destroy its integrity. Since principal p acts for $p \vee q$, a Byzantine failure of p could also compromise $p \vee q$. Conversely, the policy $o \overset{I}{\leftarrow} p \wedge q$ means that o believes that the integrity of the labeled data will be compromised only if *both* p and q fail.

3.4 Ordering labels

One label ℓ_2 is at least as *restrictive* as another label ℓ_1, written $\ell_1 \sqsubseteq \ell_2$, if it is always permissible to use information labeled ℓ_1 in a situation where information with label ℓ_2 is expected. For this to be true, ℓ_1 must offers integrity and availability guarantees at least as strong as those offered by ℓ_2. We define the *no more restrictive than* relation $\sqsubseteq$ on labels using the relations $\sqsubseteq_I$ and $\sqsubseteq_A$ on integrity and availability policies. For l_1 to be no more restrictive than l_2 in the DLM, *every* principal must believe that the integrity and availability requirements expressed by l_2 are enforced by l_1. For example, we have $\{o \overset{I}{\leftarrow} p\} \sqsubseteq \{o \overset{I}{\leftarrow} p \vee q\}$, because the left-hand label means that o believes only p has affected the data, whereas the right-hand label also permits q to affect it. Thus, the more principals have affected some information, the more restricted future use of the information becomes.

Principals are only responsible for enforcing policies that they own, but if $p_1 \succeq p_2$, then p_1 enforces all labels that p_2 owns. If labels contain multiple policies owned by multiple different principals, principals may have different views of the ordering on those labels. That is why information flow is considered acceptable only when the ordering on the labels is acceptable according to the view of *every* principal.

The formal definition of the relationship $l_1 \sqsubseteq l_2$ also takes into account the trust relationships among principals [10]. Because of trust relationships, the most restrictive integrity policy is $\bot \overset{I}{\leftarrow} \bot$, since all principals believe any principal could influence the information, and the least restrictive is $\top \overset{I}{\leftarrow} \top$, since all principals believe that only $\top$ has influenced the information (that is, it is always very trustworthy). Availability works much like integrity, with $\bot \overset{A}{\leftarrow} \bot$ meaning anyone can interfere with the labeled information's availability and $\top \overset{A}{\leftarrow} \top$ meaning that only $\top$ can stop the information from being available. A more formalized definition of $\sqsubseteq$ can be found in Appendix A. Figure 2 illustrates the lattice of the integrity and availability labels.

We define the notation $I(\ell)$ to mean the integrity policies of a label, and $A(\ell)$ to mean the availability policies of a label. The relationship $l_1 \sqsubseteq l_2$ holds exactly when the same relationship holds separately on the availability and integrity policies of the label: $\ell_1 \sqsubseteq \ell_2 \Leftrightarrow I(\ell_1) \sqsubseteq I(\ell_2) \wedge A(\ell_1) \sqsubseteq A(\ell_2)$.

3.5 Lattice operators

Since labels form a lattice, there are the usual lattice join ($\sqcup$) and meet ($\sqcap$) operators. The join of two labels gives the strongest integrity and availability that both kinds of information can flow to, and the meet gives the weakest integrity and

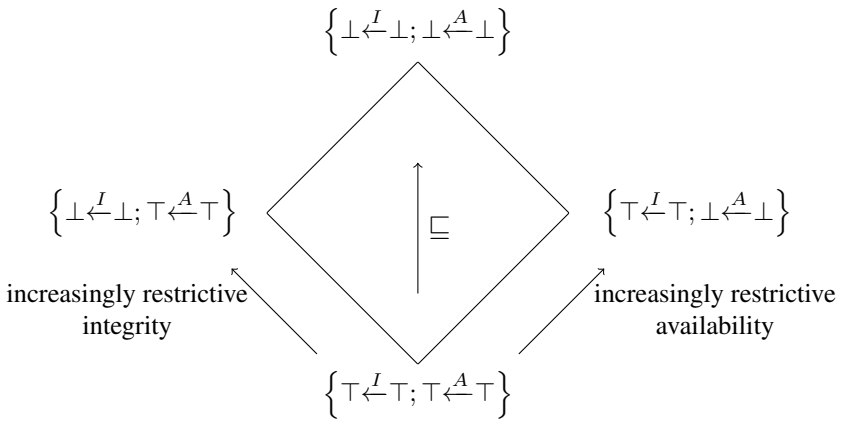
$\left\{\bot \xleftarrow{I} \bot; \bot \xleftarrow{A} \bot\right\}$
$\left\{\bot \xleftarrow{I} \bot; \top \xleftarrow{A} \top\right\}$
$\sqsubseteq$
$\left\{\top \xleftarrow{I} \top; \bot \xleftarrow{A} \bot\right\}$
increasingly restrictive integrity
increasingly restrictive availability
$\left\{\top \xleftarrow{I} \top; \top \xleftarrow{A} \top\right\}$

Figure 2. The integrity and availability lattice of labels.

availability that is allowed to flow to both kinds of information. Thus, $\sqcup$ acts as a disjunction and $\sqcap$ acts as a conjunction.

For example, $\{o \xleftarrow{I} p\} \sqcup \{o \xleftarrow{I} q\} = \{o \xleftarrow{I} p \vee q\}$, and $\{o \xleftarrow{I} p\} \sqcap \{o \xleftarrow{I} q\} = \{o \xleftarrow{I} p \wedge q\}$. A more formalized definition can be found in Appendix A.

The operations $\sqcup$ and $\sqcap$ operate separately on the integrity and availability components of labels:

$$\begin{array}{l|l} I(\ell_1 \sqcup \ell_2) = I(\ell_1) \sqcup I(\ell_2) & A(\ell_1 \sqcup \ell_2) = A(\ell_1) \sqcup A(\ell_2) \\ I(\ell_1 \sqcap \ell_2) = I(\ell_1) \sqcap I(\ell_2) & A(\ell_1 \sqcap \ell_2) = A(\ell_1) \sqcap A(\ell_2) \end{array}$$

4. Reasoning about heterogeneous trust

4.1 Threshold synthesizers

Zheng and Myers introduced the concept of *message synthesizers* in a distributed system [37, 39]. A synthesizer listens for messages from a set of hosts, and based on messages received, may produce a value with a label representing assurances that can exceed what any one message can provide.

For example, suppose a principal p awaits the receipt of messages $m_\mathsf{a}, m_\mathsf{b}, m_\mathsf{c}$ from principals a, b, and c. Let us use $\ell(m)$ to denote the label of message m. Assuming p trusts a, b, and c to send these messages, the least restrictive guarantee that can be made regarding their availability is that $A(\ell(m_\mathsf{a})) = \left\{\mathsf{p} \xleftarrow{A} \mathsf{a}\right\}$, $A(\ell(m_\mathsf{b})) = \left\{\mathsf{p} \xleftarrow{A} \mathsf{b}\right\}$, and $A(\ell(m_\mathsf{c})) = \left\{\mathsf{p} \xleftarrow{A} \mathsf{c}\right\}$. Suppose the message synthesizer $\pi_{fastest}$ listens for these messages and returns whichever message arrives first. This synthesizer produces a value that is more available (has less restrictive availability) than any one of the messages, because the value is unavailable only if all three of a, b, and c are unavailable:

$$A(\ell(\pi_{fastest})) = A(\ell(m_\mathsf{a})) \sqcap A(\ell(m_\mathsf{b})) \sqcap A(\ell(m_\mathsf{c})) = \left\{\mathsf{p} \xleftarrow{A} (\mathsf{a} \wedge \mathsf{b} \wedge \mathsf{c})\right\}$$

However, this synthesizer allows a, b, or c to affect its value, so its integrity is correspondingly lowered:

$$I(\ell(\pi_{fastest})) = I(\ell(m_\mathsf{a})) \sqcup I(\ell(m_\mathsf{b})) \sqcup I(\ell(m_\mathsf{c})) = \left\{\mathsf{p} \xleftarrow{I} (\mathsf{a} \vee \mathsf{b} \vee \mathsf{c})\right\}$$

In contrast, we might have a different message synthesizer π_{all} that returns a value only once it receives all three messages, and only if all three carry identical values. In this case, the results are available only if all the message senders are available (any one sender can render it unavailable), but the integrity of the result is much less restrictive: the result can be corrupted only if all three messages were corrupted:

$$A(\ell(\pi_{all})) = A(\ell(m_{\mathsf{a}}))\sqcup A(\ell(m_{\mathsf{b}}))\sqcup A(\ell(m_{\mathsf{c}})) = \left\{\mathsf{p} \xleftarrow{A} (\mathsf{a} \vee \mathsf{b} \vee \mathsf{c})\right\}$$

$$I(\ell(\pi_{all})) = I(\ell(m_{\mathsf{a}}))\sqcap I(\ell(m_{\mathsf{b}}))\sqcap I(\ell(m_{\mathsf{c}})) = \left\{\mathsf{p} \xleftarrow{I} (\mathsf{a} \wedge \mathsf{b} \wedge \mathsf{c})\right\}$$

The availability constraint poses an interesting problem here: sometimes we want to know if an availability constraint is met. For example, we might want to know if π_{all} received contradictory messages. We follow Zheng [37] by using the special value **none** to denote a detectably unavailable value.

A consensus protocol is a form of message synthesizer. It takes in messages, and synthesizes a consensus message, accompanied by availability and integrity guarantees that no one message could have.

4.2 Comparison to survivor sets and failure-prone sets

A different way to express the powers of an attacker is to use *failure-prone sets* of principals. [16, 17, 22]. A non-blocking protocol is one in which no failure of any subset of a failure-prone set prevents termination or progress. The complements of failure-prone sets are *survivor sets*. Every failure of any subset of a failure prone set leaves at least one survivor set failure-free.

We can therefore express a principal p_{sys} that has the least restrictive integrity or availability the system can have, as the disjunction of each of the survivor sets, each represented by the conjunction of its members.

Similarly, we can express a principal p_{attack} that has the most restrictive availability or integrity the attacker can't have, as the weakest compound principal such that it is impossible to partition the set of principals into one partition whose conjunction implies p_{sys}, and another whose conjunction implies p_{attack}. A survivor set survives iff the attacker cannot act for p_{attack}, and so in the case of crash failures, a piece of information that must be supplied by a survivor set to some principal o has the availability $o \xleftarrow{A} p_{attack}$, and in the case of Byzantine failures, the integrity $o \xleftarrow{I} p_{attack}$.

For example, in a system of four participants $P = \{\mathsf{a}, \mathsf{b}, \mathsf{c}, \mathsf{d}\}$, any one of which might fail we have:

$$\begin{aligned} p_{sys} &= (\mathsf{a} \wedge \mathsf{b} \wedge \mathsf{c}) \vee (\mathsf{a} \wedge \mathsf{b} \wedge \mathsf{d}) \vee (\mathsf{a} \wedge \mathsf{c} \wedge \mathsf{d}) \vee (\mathsf{b} \wedge \mathsf{c} \wedge \mathsf{d}) \\ p_{attack} &= (\mathsf{a} \wedge \mathsf{b}) \vee (\mathsf{a} \wedge \mathsf{c}) \vee (\mathsf{a} \wedge \mathsf{d}) \vee (\mathsf{b} \wedge \mathsf{c}) \vee (\mathsf{b} \wedge \mathsf{d}) \vee (\mathsf{c} \wedge \mathsf{d}) \end{aligned}$$

Using this construction, any protocol phrased in terms of survivor and failure-prone sets can be expressed in terms of labels instead. Since the condition that no more than f failures occur can be converted to survivor and failure-prone sets, it can also be expressed in terms of labels.

There is also a certain kind of equivalence between labels and survivor sets: any label can be expressed in terms of (for each principal) two collections of survivor and failure-prone sets, describing the crash and Byzantine failures that principal tolerates.

5. Heterogeneous fast consensus

As an example of a protocol adapted for a heterogeneous trust environment, we present Heterogeneous Fast Consensus, a generalization of the Bosco Fast Consensus protocol [32]. Fast Consensus is a one-round protocol that can, in the best case, decide on a consensus value in one communication step. If it fails to decide, some underlying consensus is used; this can be another round of Fast Consensus. The desirable properties of a traditional consensus protocol [6, 32] can be generalized for heterogeneous trust:

Agreement: If two gurus decide, then they decide the same value. Also, if a guru decides more than once, it decides the same value each time.

Unanimity: If all correct participants have the same initial value v, a guru that decides must decide v.

Validity: If a participant decides a value, and all participants are correct, then that value was proposed by some participant.

Progress: Under the assumption that underlying consensus terminates, all gurus eventually decide.

Termination: Under the assumption that the network delivers concurrent messages in random order, and underlying consensus is simply a recursive invocation of another instance of Fast Consensus, all gurus eventually decide with probability 1.

Algorithm 1 contains the pseudo-code for Heterogeneous Fast Consensus, and Appendix B contains the proof of correctness. Each participant broadcasts its starting value to each participant (including itself). Once it can no longer be sure it will receive any more messages (given its failure assumptions), the participant looks over the values it has received. If a set of identically-valued messages meets the high, *decision*, threshold, that participant can decide that value. What a participant does to *decide* varies depending upon *underlying-consensus*(), but if *underlying-consensus*() is *fast-consensus*(), it simply broadcasts the decided value one last time:

$decide(v)$: {
    send message with value v **to** each participant;
    return v
}

If some value has a set of messages which meet a lower *change* threshold, then the participant enters *underlying-consensus*() with that as its starting value. Otherwise, the participant picks a value v from those received using some *selection-function*(), and invokes *underlying-consensus*(v).

The function *selection-function*() varies depending upon the desired properties of the protocol and on *underlying-consensus*(). If *underlying-consensus*() is *fast-consensus*(), the protocol may converge fastest if *selection-function* always selects the first in some arbitrary but consistent ordering of the input values. However,

```
1  Function fast-consensus(v_p):
2      send message with value v_p to each participant;
3      Upon receipt of a message:
4          R ⟵ the set of messages received thus far;
5          if sufficiently-available(R) then
6              forall the unique values v in m ∈ R do
7                  S ← {m|(value of m) = v ∧ m ∈ R};
8                  if sufficient-to-decide(S) then
9                      return decide(v);
10                 else if sufficient-to-change(S) then
11                     return underlying-consensus(v)
12                 end
13             end
14             return underlying-consensus(
                                  selection-function(R))
15         end
16     end
17 end
```

Algorithm 1: Pseudo-code for Fast Consensus [32]. The input is the starting value for a given participant. Functions *sufficiently-available*(), *sufficient-to-decide*, and *selection-function* are discussed elsewhere. Function *underlying-consensus* is the consensus protocol to be used in the case that this one-round consensus fails to decide. It is assumed to take as input a participant's starting value, and it can be another round of *fast-consensus*(). This pseudo-code assumes a language mechanism for generating a list of expected messages for this particular consensus round.

this deterministic strategy may permit a Byzantine attacker to prevent agreement in each round. It is therefore prudent to incorporate some randomness in *selection-function*(), as in RS-Bosco [32].

5.1 Modeling simple homogeneous trust

In the traditional case, in which all n participants believe that any c participants may crash, and any b participants may fail in Byzantine fashion (which traditionally includes crashing, so $c \geq b$), the thresholds are fairly straightforward [32]:

- *sufficiently-available*$(R) \equiv |R| \geq n - c$
- *sufficient-to-decide*$(S) \equiv |S| > \frac{n+c}{2} + b$
- *sufficient-to-change*$(S) \equiv |S| > \frac{n-c}{2}$

Requirements: This protocol has similarly straightforward requirements on the values of n, c and b. A participant can only be sure of receiving $n - c$ votes, and needs more than $\frac{n+c}{2} + b$ to decide anything, so it is required that $n > 3c + 2b$. If only crash failures are expected ($b = 0$), then this means $n > 3c$, and if only Byzantine failures are expected ($b = c$), then this means $n > 5b$. These bounds are known to be tight for single-communication-step consensus [32].

5.2 Heterogeneous Trust

In an information flow setting, labels are assigned to each piece of information to describe properties such as availability and integrity, characterizing who might make that information unavailable or affect its contents (section 3). In this case, each message m has a corresponding label $\ell(m)$. A set of messages taken together can be used to synthesize a value with more integrity or availability than any message alone (section 4.1).

Each participant may have a different idea of the possible failures. A participant p can phrase its requirement for when it believes it cannot be certain of receiving any more messages as, for some availability policy A^p_{sys} unique to p, the condition $\bigsqcap_{m \in R} A(\ell(m)) \sqsubseteq A^p_{\text{sys}}$. Label A^p_{sys} may be constructed as in section 4.2. As A^p_{sys} should contain only policies owned by p, and the only portions of the label on m which p must enforce are those owned by p, we can define:

$A_{\text{sys}} = \bigsqcap_{p \in P} A^p_{\text{sys}}$, and can then write:

$$\textit{sufficiently-available}(R) \equiv \bigsqcap_{m \in R} A(\ell(m)) \sqsubseteq A_{\text{sys}}$$

Similarly, a threshold C^p, composed of both availability and integrity policies, represents when a participant feels a set of messages should be sufficient to force it to change its starting value to a particular value in *underlying-consensus*. We define $C = \bigsqcap_{p \in P} C^p$, and:

$$\textit{sufficient-to-change}(S) \equiv \bigsqcap_{m \in S} \ell(m) \sqsubseteq C$$

Similarly, a threshold D^p, composed of both availability and integrity policies, represents when a participant feels a set of messages should be sufficient to decide on a value carried by all of those messages. We define $D = \bigsqcap_{p \in P} D^p$, and:

$$\textit{sufficient-to-decide}(S) \equiv \bigsqcap_{m \in S} \ell(m) \sqsubseteq D$$

5.3 Requirements

As in the case of simple homogeneous trust, heterogeneous trust imposes requirements on the values of A_{attack}, A_{sys}, I_{attack}, I_{sys}, C, and D, as well as on the labels of messages sent as part of the protocol.

Let m^p_q designate a message from participant p to participant q. Let A^p_{attack} designate the most restrictive availability label participant p believes an attacker can't block. Let I^p_{attack} designate the most restrictive integrity label participant p believes an attacker can't influence.

From the perspective of a principal p, therefore, we can define failure-prone sets for Byzantine and crash failures, called Liar sets and Crash sets:

$$L \text{ is a liar set } \Leftrightarrow \left(\bigsqcap_{p \in P, q \in L} \left\{ p \xleftarrow{I} q \right\} \right) \not\sqsubseteq I^p_{\text{attack}}$$

$$H \text{ is a crash set } \Leftrightarrow \left(\bigsqcap_{p \in P, q \in H} \left\{ p \xleftarrow{A} q \right\} \right) \not\sqsubseteq A^p_{\text{attack}}$$

It is possible, for example, to construct A^p_{sys} given A^p_{attack} from survivor sets: the crash sets' compliments. This notation simplifies future expressions.

A participant may not be able to send a message if it is itself unavailable. Therefore, from the perspective of participant q, the availability of a message m^p_q from participant p is limited by this constraint:

$$\left\{q \xleftarrow{A} p\right\} \sqsubseteq A\big(\ell\big(m^p_q\big)\big) \tag{1}$$

Each participant collects received messages R until the meet of their labels of all received messages is no more restrictive than A_{sys}. This places a limit on *viable* systems. For no participant should it ever be the case that the attacker (as perceived by that participant) can prevent such a set of messages from arriving:

$$\forall q \in P.\, \forall F \subseteq P.\, (q \notin F) \wedge \left(\bigsqcap_{f \in F} \left\{q \xleftarrow{A} f\right\} \not\sqsubseteq A_{\text{attack}} \right) \Rightarrow \bigsqcap_{p \in (P-F)} \big(A\big(\ell\big(m^p_q\big)\big)\big) \sqsubseteq A_{\text{sys}} \tag{2}$$

If *underlying-consensus*() is *fast-consensus*(), no participant should progress to the next round unless it can complete that round with certainty. This condition implies that if messages from a set of participants messages can propel p to the next round, but not q, then p must be able to progress with that set, without q.

$$\begin{gathered} \forall r, q \in P.\, \forall R \subseteq P. \\ \left(\left(\bigsqcap_{p \in S} A(\ell(m^p_r)) \sqsubseteq A_{\text{sys}} \right) \wedge \left(\bigsqcap_{p \in S} A(\ell(m^p_q)) \not\sqsubseteq A_{\text{sys}} \right) \right) \\ \Downarrow \\ \left(\bigsqcap_{p \in S - \{q\}} A(\ell(m^p_r)) \sqsubseteq A_{\text{sys}} \right) \end{gathered} \tag{3}$$

Another requirement is that no set of failures a participant believes might happen should prevent that participant from deciding. Additionally, any set of messages that makes a correct participant decide should also be sufficient to dictate that participant's value in *underlying-consensus*.

$$\left(A_{\text{sys}} \sqcap \left\{ \top \xleftarrow{I} \top \right\} \right) \sqsubseteq D \sqsubseteq C \tag{4}$$

If a participant p decides, it is useful to talk about *decider* sets and their complements, *wrong sets*. A decider set is any set of principals whose messages can make p decide. In much the same way that we might construct A_{attack} and *crash sets* from A_{sys}, we can construct W and *wrong sets* from D.

$$E \text{ is a decider set } \Leftrightarrow \bigsqcap_{e \in E} \ell\big(m^e_p\big) \sqsubseteq D$$

$$G \text{ is a wrong set } \Leftrightarrow \bigsqcap_{g \in G} \ell\big(m^g_p\big) \not\sqsubseteq W$$

A wrong set is any set of participants who broadcast, or may have broadcast, a value other than the one p decided during the round in which p decided. The label W dictates the least restrictive availability and integrity that no wrong set may have.

Assuming *underlying-consensus* has unanimity, Fast Consensus assures agreement by ensuring that if one guru decides, then all correct participants who have

$$
\begin{array}{rlcl}
 & A^{\mathsf{a}}_{\text{sys}} & = & \left\{\mathsf{a} \xleftarrow{A} (\mathsf{a} \wedge \mathsf{c} \wedge \mathsf{d} \wedge \mathsf{e})\right\} \\
\forall p \in \{\mathsf{b},\mathsf{c},\mathsf{d},\mathsf{e}\}. & A^{p}_{\text{sys}} & = & \left\{p \xleftarrow{I} p \wedge \left(\bigvee_{x,y,z \in \{\mathsf{b},\mathsf{c},\mathsf{d},\mathsf{e}\}} (x \wedge y \wedge z)\right)\right\} \\
\forall p \in P. & C^p & = & \left\{p \xleftarrow{I} \bigvee_{x,y \in \{\mathsf{b},\mathsf{c},\mathsf{d},\mathsf{e}\}} (x \wedge y)\right\} \\
\forall p \in P. & D^p & = & \left\{p \xleftarrow{I} \bigvee_{x,y,z \in \{\mathsf{b},\mathsf{c},\mathsf{d},\mathsf{e}\}} (x \wedge y \wedge z)\right\} \\
\forall p,q \in P. & A\big(\ell\big(m^{\mathsf{q}}_p\big)\big) & = & \left\{p \xleftarrow{A} q\right\} \\
\forall p,q \in P. & I\big(\ell\big(m^{\mathsf{q}}_p\big)\big) & = & \left\{p \xleftarrow{I} q\right\}
\end{array}
$$

Figure 3. Using *fast-consensus* as *underlying-consensus*, and a *selection-function* of their choice (e.g., random selection), the participants in the example can execute fast consensus with these threshold values and labels.

not yet decided enter *underlying-consensus* with the value decided. Fast Consensus therefore has some requirements pertaining to C and D.

If p is a guru, then none of p's perceived wrong sets, combined with any of p's perceived liar sets, should be able to change anyone's vote. Otherwise, a set of liars combined with a set of participants p is aware may have broadcast something other than it decided could prevent correct participants from entering *underlying-consensus* with unanimous values.

$$
\forall p,q.\left(\left(\bigsqcap_{l \in L}\left\{p \xleftarrow{I} l\right\}\right) \not\sqsubseteq I^p_{\text{attack}} \wedge \left(\bigsqcap_{h \in H}\left\{p \xleftarrow{A} h\right\}\right) \not\sqsubseteq W\right) \Rightarrow \left(\bigsqcap_{x \in L \cup H} \ell\big(m^x_q\big)\right) \not\sqsubseteq C \tag{5}
$$

We also require that if one guru decides, it must be impossible for any other correct participant not to enter *underlying-consensus* with the decided value. This means that for any group of liars L, and any group of crashers H, and any additional group J, if messages from $L \cup H \cup J$ make a guru decide, then messages from J, being the only ones guaranteed to get through to the other participants, must make those participants enter *underlying-consensus* with the decided value.

$$
\begin{array}{c}
\forall p,q \in P. \forall L,H,J \subseteq P. \\
\left(\left(\left(\bigsqcap_{l \in L}\left\{p \xleftarrow{I} l\right\}\right) \not\sqsubseteq I^p_{\text{attack}}\right) \wedge \left(\left(\bigsqcap_{h \in H}\left\{q \xleftarrow{A} h\right\}\right) \not\sqsubseteq A^q_{\text{attack}}\right) \wedge \left(\left(\bigsqcap_{x \in L \cup H \cup J} \ell\big(m^x_p\big)\right) \sqsubseteq D\right)\right) \\
\Downarrow \\
\left(\bigsqcap_{j \in J} \ell\Big(m^j_q\Big)\right) \sqsubseteq C
\end{array} \tag{6}
$$

5.4 Example

Returning to the example introduced in section 1, we can now synthesize a consensus protocol for Alice, Bob, Carol, Dave, and Eve. For brevity, the letters a–e are used to represent the five participants.

Alice believes that Bob may fail in a Byzantine fashion (lose integrity). She does not believe any other failures may occur.

Bob, Carol, Dave, and Eve each believe Alice can fail in a Byzantine fashion (lose integrity), and believe that at most one other participant may crash (lose availability) as well.

These trust assumptions are captured by the following labels:

$$\begin{aligned} A^{\mathsf{a}}_{\text{attack}} &= \left\{ \mathsf{a} \xleftarrow{A} \mathsf{a} \vee \mathsf{c} \vee \mathsf{d} \vee \mathsf{e} \right\} \\ I^{\mathsf{a}}_{\text{attack}} &= \left\{ \mathsf{a} \xleftarrow{I} \mathsf{a} \vee \mathsf{c} \vee \mathsf{d} \vee \mathsf{e} \right\} \\ \forall p \in \{\mathsf{b},\mathsf{c},\mathsf{d},\mathsf{e}\}. A^{p}_{\text{attack}} &= \left\{ p \xleftarrow{A} p \vee \left(\bigvee_{q,r \in \{\mathsf{b},\mathsf{c},\mathsf{d},\mathsf{e}\}} (q \wedge r) \right) \right\} \\ \forall p \in \{\mathsf{b},\mathsf{c},\mathsf{d},\mathsf{e}\}. I^{p}_{\text{attack}} &= \left\{ p \xleftarrow{I} \mathsf{b} \vee \mathsf{c} \vee \mathsf{d} \vee \mathsf{e} \right\} \end{aligned}$$

Solution: Search of the space of threshold labels (A_{sys}, C, D) reveals that there are indeed thresholds meeting the participants' requirements, as well as the requirements of Fast Consensus. Using *fast-consensus* as *underlying-consensus*, and a *selection-function* of their choice (from a theoretical perspective, random selection works), they can execute Fast Consensus with the threshold values and labels found in Figure 3.

One counterintuitive insight provided by our analysis is that there are occasions in which Alice listens to Bob, despite the fact that she does not trust him at all. From the label analysis, this falls out from C^{a} and D^{a}, defined in Figure 3, which can be shown to satisfy the threshold requirements.

Intuitively, the logic is this: If Alice is correct, and Bob is Byzantine, and so everyone else is a chump, then it doesn't matter what Alice decides, so long as she does decide. If, on the other hand, Bob is only crash-failure prone, then Alice can't decide having heard only from, for example, Carol and Dave, because Carol and Dave may have heard two votes from Eve and Bob, different from what they voted for, and change their votes next round. As a result, Alice would have decided something different from what the others decide, despite no one having been wrong (or failing). Therefore, Alice must also wait to hear from Bob or Eve, to ensure that in the event that Bob is only crash-prone, Carol, Dave, and Eve will decide the same as what Alice decides.

Evaluation: In this trust configuration, Eve tolerates the Byzantine failure of Alice and the simultaneous crash failure of Bob. Therefore, traditional (homogeneous) Bosco would have to tolerate one Byzantine failure and one *additional* crash failure, requiring a total of 9 participants. Already Heterogeneous Fast Consensus has a clear advantage for this scenario: it requires only 5 participants to tolerate this trust configuration, whereas traditional Bosco requires recruiting at least 4 more trustworthy participants, who also slow the system down. We simulated 1000 instances in which 9 participants participated in Bosco tolerating 1 Byzantine failure and 2 total crash failures. The network delivered messages in each round in an order drawn uniformly at random from all possible orderings. The *selection-function*() used chose a value uniformly at random from the set of messages received.

We also ran the same simulation for our Heterogeneous Fast Consensus implementation, and for each calculated the mean (with standard error) probability of a participant deciding after each round. All participants began with different values.

For the case in which no failures occurred, the results are in figure 4. Not only did Heterogeneous Fast Consensus converge quickly, in a median of 3 rounds,

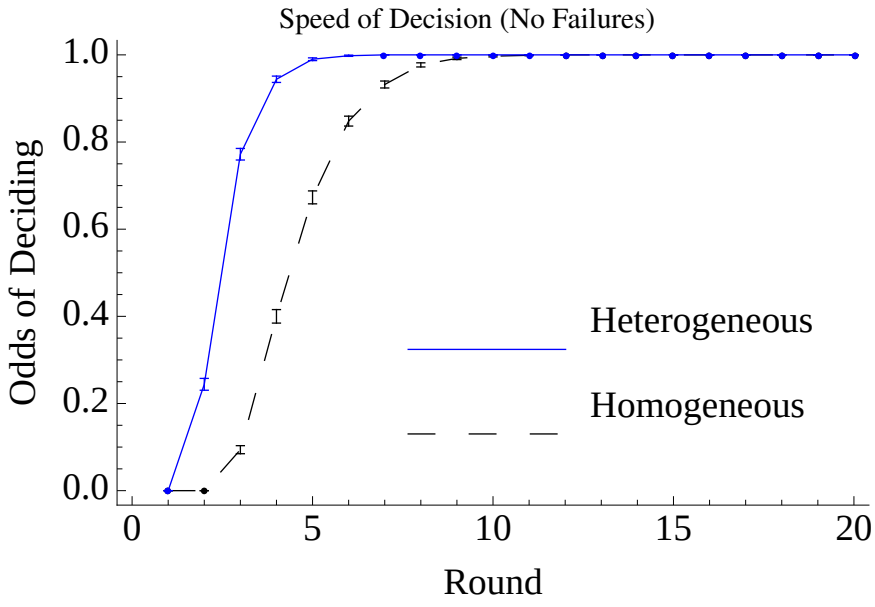

Figure 4. The probability of a participant deciding in each round, mean over 1000 samples, in our 5-participant Heterogeneous Fast Consensus protocol, and a traditional 9-participant Bosco protocol. Standard error bars shown.

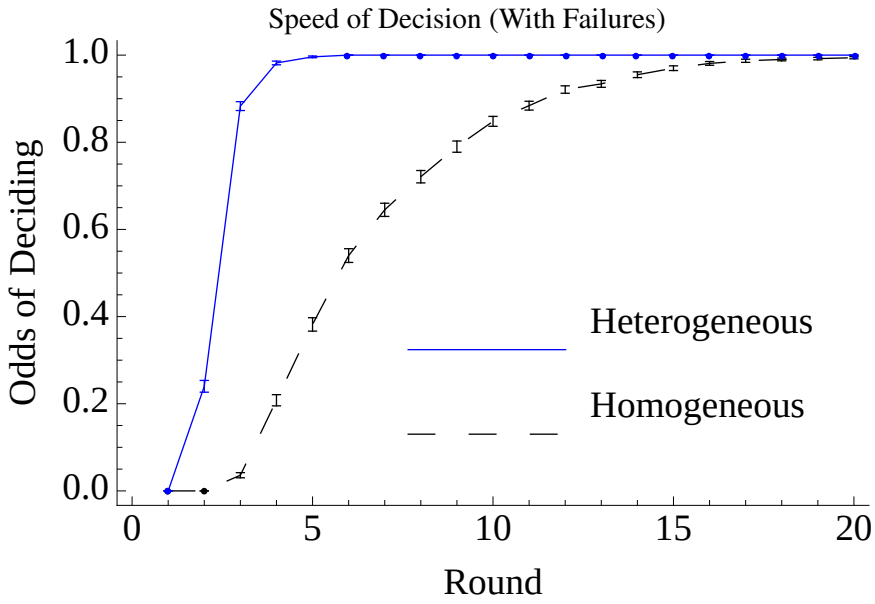

Figure 5. The probability of a participant deciding in each round, mean over 1000 samples, in our 5-participant Heterogeneous Fast Consensus protocol, and a traditional 9-participant Bosco protocol. In this case, Alice is Byzantine, and Bob has crashed. Standard error bars shown.

but it converged much faster than Homogeneous Bosco, in a median time of 5 rounds. The gap was even wider in the 95th percentile, where Heterogeneous Fast Consensus took 5 rounds, and Homogeneous Bosco took 8.

For the case in which Alice has failed in a Byzantine fashion (specifically, she proposes a new, never-before-seen value each round), and Bob has crashed, the difference is even greater. With the reduced contention of fewer active participants, Heterogeneous Fast Consensus converges even faster, the 95th percentile deciding by round 4, while the homogeneous case takes until round 6 for the median to decide, and round 12 before the 95th percentile decided. The full results are in figure 5.

These results suggest that customizing the protocol to the heterogeneous trust configuration yields clear advantages in both resource requirements and speed.

Determining threshold labels

One challenge of the Heterogeneous Fast Consensus protocol is finding appropriate threshold labels to satisfy the requirements of section 5.3. This is an offline computation, so performance is not critical. We expressed the requirements using quantifier-free bitvector logic (QF_BV), and used the Z3 SMT solver to find solutions. Generating four- to six-participant protocols took a few minutes for each protocol. The time to generate protocols increases as the number of participants increases, but one does not normally generate consensus protocols for very large numbers of participants in any case.

We have made the search script available at `https://www.dropbox.com/s/akv957fmrqsn803/pysmt.zip?dl=0`.

Generalized heterogeneous fault-tolerant protocols will have, in general, parameters specific to the needs of their participants. For example, heterogeneous fast consensus has the thresholds A_{sys}, C, and D, which must be crafted to fit the requirements of the protocol and the specific distrust of the participants. The exact nature of the complexity of such a search is unclear. For any given set of labels, it is easy (a polynomial-time computation) to check that they meet these constraints. The problem is therefore in NP, but the precise hardness of the problem remains future work.

5.5 Very Fast Consensus

This construction of fast consensus suggests another way to exploit heterogeneous trust. Suppose a client wishes to submit a request to a group of servers, from which it requires an answer upon which the servers must reach consensus. Suppose this operation is extremely latency-sensitive, and the goal is best-case performance of one communication to the servers, no communication time between servers, and one communication to the client. This construction of fast consensus extends to cover this extremely latency-sensitive case simply by adding the client to the set of participants, as a participant whom the servers don't trust at all. The client thus receives messages as a participant, and can set its own thresholds for when it is satisfied consensus has been reached (subject, of course, to the protocol's requirements). This is not efficient in terms of bandwidth, but highly efficient in terms of latency, as the client need not wait for any communication between servers before servers send it a response.

6. OARcast and Nysiad

To further illustrate the utility of our approach, we present a second generalization of a distributed protocol, using the same techniques based on information flow analysis. In this case, we generalize Nysiad, an algorithm for converting crash-tolerant state-machine systems ([30]) to Byzantine-tolerant ones [13, 14]. In particular, the conversion process allows a faulty "sender" participant to interfere with the availability, but not the integrity, of a message.

6.1 Ordered Asynchronous Reliable Broadcast

At its core, Nysiad is built around Ordered Asynchronous Reliable broadcast (OARcast), a protocol that is useful in its own right. OARcast, as presented in [13, 14], works as follows.

There exists a set of participants known as *echoers*. One special echoer, the *designated sender*, may wish to broadcast a message m such that all other echoers receive it. The goal of the protocol is that all messages broadcast should arrive in the same order at all nodes: the order in which they were sent.

Each message is assumed to be signed by both its author and its sender (which may not be the author if the message is relayed through an echoer). Each message is also assumed to contain a sequence number, assigned sequentially by the author.

Each echoer will echo any new message (that it's not seen before) from the designated sender to all other echoers. Any echoer that receives two messages signed by the designated sender, and containing the same sequence number, that have different values, ceases operation.

An echoer *delivers* a message (that is, produces a value) when it has delivered messages with all lesser sequence numbers, and the set of identical messages its received for this sequence number meets a condition. This condition can be expressed with label threshold synthesizers, instead of waiting for a specific number of identical messages (See section 4.1). The basic requirement, that no set of messages from echoers should allow two guru echoers to deliver different messages, remains the same. Specifically, if each echoer e has an integrity value I_a^{e}, defined to be strictly less restrictive than the integrity of any attacker it tolerates, (expressed $\sqsubset$, meaning $x \sqsubset y \Leftrightarrow x \sqsubseteq y \wedge y \not\sqsubseteq x$), and has some integrity threshold T_{e} for delivering a message, the following condition must hold to ensure gurus never deliver different messages:

$$\begin{gathered}\forall \mathsf{e}, \mathsf{e}' \in P. \forall A, B, C \subseteq P : B \cap C = \{\}. \\ \left(I_a^{\mathsf{e}} \sqsubset \bigsqcap_{\mathsf{q} \in A} I\big(\ell\big(m_{\mathsf{e}}^{\mathsf{q}}\big)\big)\right) \wedge \left(\bigsqcap_{\mathsf{q} \in (A \cup B)} I\big(\ell\big(m_{\mathsf{e}}^{\mathsf{q}}\big)\big) \sqsubseteq T_{\mathsf{e}}\right) \\ \Downarrow \\ \left(\bigsqcap_{\mathsf{q} \in (A \cup C)} I\big(\ell\big(m_{\mathsf{e}'}^{\mathsf{q}}\big)\big) \not\sqsubseteq T_{\mathsf{e}'}\right)\end{gathered} \tag{7}$$

The proof of correctness proceeds exactly as in [13, 14]. Concerning safety properties, it should be clear that:

- No two gurus can deliver different values for the same message sequence number.
- If the designated sender is correct, then no correct participant (even a chump) can simulate delivery of messages in any order other than that provided by the designated sender.

For liveness, however, we need a guarantee that:

- A guru will always deliver any message sent by a correctly functioning sender, so long as no attacker can compromise a set of echoers such that the remainder is insufficient for the guru. Formally:

$$\forall \mathsf{e} \in P. \forall A \cup B = P. I_a^{\mathsf{e}} \sqsubset \bigsqcap_{\mathsf{q} \in A} I(\ell(m_{\mathsf{e}}^{\mathsf{q}})) \Rightarrow \bigsqcap_{\mathsf{q} \in (B)} I(\ell(m_{\mathsf{e}}^{\mathsf{q}})) \sqsubseteq T_e$$

With this additional requirement, no attacker can prevent a guru from reaching its delivery conditions.

The integrity of the delivered message cannot therefore only be interfered with by the echoer's perceived attacker.[2] In other words, the integrity of the delivered message is I_a^{e}. Availability is limited by the sender, as well as any conditions under which integrity is violated, calculated as follows:

$$A_{\mathsf{s}} \sqcap \bigsqcup_{M \subseteq \left\{ m_{\mathsf{e}}^{\mathsf{p}} | p \in P \wedge \sqcap_{m \in M} I(\ell(m)) \sqsubseteq I_a^{\mathsf{e}} \right\}} \left(\bigsqcap_{m \in M} A(\ell(m)) \right)$$

where A_{s} is the availability of the designated sender.

6.2 Nysiad

Nysiad, a translation mechanism from arbitrary crash-tolerant protocols to Byzantine tolerant ones, also presented in [13, 14], can be performed using the heterogeneous OARCast.

The idea of Nysiad is to take any crash-tolerant system consisting of a collection of deterministic state machines, and simulate it on each of a group of participants. For each state machine in the original system, the participants form an OARcast, and one participant is designated as the sender for that machine. Each participant simulates the delivery of a message to each simulated state machine only if it has itself calculated an identical simulated message sent to that simulated machine, and it has received an identical message via that machine's OARcast. In this way, all gurus simulate identical executions, with identical message delivery ordering. OARcast ensures that a failed participant cannot force two gurus to perceive message delivery in different orders, and the requirement that each participant derive the simulated messages themselves ensures that no failed participant can force another participant to deliver an incorrect value. The availability of simulated messages is thus limited by the availability of the simulated sender machine's designated sender participant, as well as the availability of the sender machine's OARcast. The integrity of simulated messages (which really is all in the ordering) is limited only by the integrity of the sender machine's OARcast, and not the designated sender participant. Additionally, all information simulated on a given participant is limited by that participant's availability and integrity.

It is notable that, in the f-failure tolerant case, a Nysiad conversion of a $3f + 1$ crash tolerant Bosco instance (with a deterministic selection function) is a $3f + 1$ Byzantine tolerant consensus protocol with best-case two message sends from proposal to decision, putting it on par (by this very simple metric), with Fast

[2] Integrity, as far as OARcast is concerned, is a property of uniform and guaranteed message delivery, and not content, which may have additional constraints.

Byzantine Paxos [8, 19]. A generalized version of Heterogeneous Fast Consensus can be likewise constructed for specific use cases.

As well as providing insightful cases in heterogeneous trust reasoning, the Nysiad algorithm, already a useful tool in constructing Byzantine-tolerant protocols, generalizes into a useful tool in heterogeneous trust based algorithms.

7. Future work

The tools, techniques, and examples in this work are meant to provide a framework for reasoning about and constructing fault-tolerant distributed protocols. We hope that protocol designers will expand on this approach to to develop novel protocols. Heterogeneous Fast Consensus is both a novel protocol and a useful example of applying information-flow techniques to fault tolerance. More efficient methods for synthesizing threshold labels remain desirable.

The holy grail for heterogeneous trust would be a procedure for transforming any existing fault-tolerant protocol into a generalized version that exploits heterogeneous trust. Such a procedure would require a way to automatically derive necessary requirements on the trust configuration and the protocol instances. While it is clear that such a procedure will not always be computable (a desirable property might be termination, and it is impossible to compute requirements for termination in general [5]), it may be feasible for useful cases.

The heterogeneous trust model of failure is extremely rich, but it does not take into account notions of self-interest, and so there is room for complementary work integrating game theory and selfish participants into this richer space [1, 3]. One might envision, for example, protocols in which participants can derive the specific implementations of an algorithm in which they take part (as opposed to centrally determining this beforehand) using knowledge of the trust configurations, and the belief that others will derive their implementations selfishly.

Finally, we have deliberately ignored confidentiality in this work, but confidentiality is also conducive to analysis using static information-flow methods [28]. Taking confidentiality into account is likely to add additional constraints to protocol design.

8. Conclusion

In our increasingly complex, interconnected world, under varied and changing threats and system models, it is critical to design systems that can operate in environments where participants make differing trust assumptions about the availability and integrity of information and of other participants. We propose the use of information-flow labels describing integrity and availability as a way to express those requirements and situations in a general manner, and to provide a rigorous framework for reasoning about protocols using heterogeneous trust. Our generalization of the Bosco Fast Consensus protocol [32], developed with this methodology, is capable of tolerating trust configurations for which traditional fast consensus fails, or would be dramatically less efficient. Properties such as Agreement, Unanimity, Validity, and Termination can be generalized for the heterogeneous case, in which some but not all correct participants make incorrect assumptions about failure.

Likewise, our generalizations of OARcast and Nysiad [13, 14], and even basic message synthesizers may serve as useful tools and building blocks in the de-

velopment of future protocols that use heterogeneous trust. The analysis of these example algorithms should serve to help others gain insight in future endeavors.

We expect that our new approach will be useful for generalizing other fault tolerant protocols to a heterogeneous trust environment, and we hope it will lead to more efficient ways to build trustworthy systems.

$$\frac{\forall\{p \in P\}.\forall X \in \{A, I\}.\left(\left(\wedge\left\{a \,\middle|\, (p \succeq o) \wedge \left(o \xleftarrow{X} a\right) \in \ell_1\right\}\right) \succeq \left(\wedge\left\{a \,\middle|\, (p \succeq o) \wedge \left(o \xleftarrow{X} a\right) \in \ell_2\right\}\right)\right)}{\ell_1 \sqsubseteq \ell_2}$$

Figure 6. Ordering on labels.

A. Principal and Lattice Formalisms

A.1 Compound Principals and $\succeq$

The rules for reasoning about compound principals:

$$(p_1 \wedge p_2) \succeq p_1 \quad \Big| \quad p_1 \succeq (p_1 \vee p_2) \quad \Big| \quad \frac{p_1 \succeq p_2 \quad p_2 \succeq p_3}{p_1 \succeq p_3} \quad \Big| \quad \frac{p_1 \succeq p_3 \quad p_2 \succeq p_3}{(p_1 \vee p_2) \succeq p_3} \quad \Big| \quad \frac{p_1 \succeq p_2 \quad p_1 \succeq p_3}{p_1 \succeq (p_2 \wedge p_3)}$$

A.2 Labels

Labels are sets of policies.

A.2.1 $\sqcup$ and $\sqcap$

For availability or integrity labels:

$$\begin{aligned} I(\ell_1 \sqcup \ell_2) &= \left\{(u_1 \vee u_2) \xleftarrow{I} (p_1 \vee p_2) \,\middle|\, u_1 \xleftarrow{I} p_1 \in I(\ell_1) \wedge u_2 \xleftarrow{I} p_2 \in I(\ell_2)\right\} \\ A(\ell_1 \sqcup \ell_2) &= \left\{(u_1 \vee u_2) \xleftarrow{A} (p_1 \vee p_2) \,\middle|\, u_1 \xleftarrow{A} p_1 \in A(\ell_1) \wedge u_2 \xleftarrow{A} p_2 \in A(\ell_2)\right\} \\ \ell_1 \sqcap \ell_2 &= \ell_1 \cup \ell_2 \end{aligned}$$

A.2.2 The lattice ordering

$\sqsubseteq$ on labels is intuitively defined in Section 3.4. Formally, it is defined in Figure 6. This definition is similar to Stephen Chong's [9].

A.2.3 Equality

Labels ℓ_1 and ℓ_2 are considered *equal*, or at least to lie in the same equivalence class, iff $(\ell_1 \sqsubseteq \ell_2) \wedge (\ell_2 \sqsubseteq \ell_1)$.

B. Heterogeneous consensus proofs

B.1 Agreement

If two gurus decide, they can do so either in the same round (of fast consensus), or one can decide in fast consensus, and the other in *underlying-consensus*.

B.1.1 Same Round:

No two gurus decide different values in the same round, by (5), (4), and $W \sqsubseteq A_{\text{attack}} \sqcap I_{attack}$ (from the definition of W). In particular, in order for a participant to decide, a group of participants must send messages with a meet featuring a label $\sqsubseteq D$. Therefore, the meet of labels of messages from participants who either lied to or did not send the decided value to the first participant is $\not\sqsubseteq C$. Given that $D \sqsubseteq C$, no message synthesized from participants who either lied to or did not send the decided value to the first participant is $\sqsubseteq D$. Therefore, no participant can decide any value other than the one decided by the first participant in the same round. Therefore, if two participants decide in the same round, they decide the same value.

B.1.2 Different Rounds:

By (6), if one guru decides, then for each other correct participant, there exists—among the participants from whom the guru has received messages—some subset J that is correct and whose messages are received by the other participant. Furthermore, J is sufficient to change the vote of that other participant to the value that has been decided.

Therefore, all correct participants (who have not yet decided) enter *underlying-consensus* with the decided value as their starting value. Assuming the first deciding participant behaves at least as a participant in *underlying-consensus* (as is the case for the given decide procedure when *underlying-consensus* is fast consensus), then if *underlying-consensus* guarantees unanimity, then all gurus will decide the same value as the first participant, and so all future decisions by gurus will agree with the first participant.

This protocol does not permit the possibility of the same guru deciding twice in a round of fast consensus, and so agreement of *underlying-consensus* combined with unanimity of *underlying-consensus* guarantees that any guru who decides twice must decide the same value both times.

B.2 Unanimity

Given (4), if all correct participants send messages of the same value, the meet of the labels of those messages is $\sqsubseteq D$. By (6), this requires that even in the presence of attackers, all correct participants will receive a set of messages with the "correct" value such that the meet of their labels $\sqsubseteq C$. All of the correct participants will therefore hold the same value when moving into *underlying-consensus*.

Therefore, if *underlying-consensus* has Unanimity, then so does Fast-consensus.

If *underlying-consensus* is fast-consensus, then no correct participant can decide any value other than the correct value. Given that in each round, all correct participants will enter with the same value, guaranteeing they do so in the next round, no correct participant will ever broadcast any other value. It is possible for a correct participant to receive all the messages from the correct participants first, and therefore decide on the correct value.

B.3 Validity

The decision procedure of this protocol only allows a correct participant to decide on an element of the set of values from received messages. Because received messages must have been sent (network assumption), we have validity.

B.4 Progress

From the perspective of any guru:

Given (1) , (2), and (3), any attacker with availability $\not\sqsubseteq A_{\text{attack}}$ would be unable to violate the system availability assumptions $\left(A_{\text{sys}}\right)$ of any set of participants such that the meet of the labels of the messages of the remainder $\not\sqsubseteq A_{\text{sys}}$.

Therefore, any guru can always expect a set of messages such that the meet of the availability of their labels $\sqsubseteq A_{\text{sys}}$, and can therefore always either decide or move on to *underlying-consensus*.

B.5 Termination

If p is a guru, then under the "random network" assumption, with some non-zero probability p will in some round get messages from all correct processes, the meet of the labels of which are $\sqsubseteq A_{\text{sys}}$, and so it will move on the next round.

There is likewise some non-zero probability that all correct participants will receive messages in the same order as p.

From requirement 3, the availability of the meet of the labels of the set of messages p received must be enough to carry a set of participants into the next round that will allow p to make further progress. (The combined availability of their messages to p must be $\sqsubseteq A_{\text{sys}}$.)

From the structure of the protocol, and requirements 5 and 6, no two correct processes should be forced to select different values after having received identical sets of messages.

Therefore, if *selection-function* has a non-zero probability of selecting each item in the input set, there is a non-zero probability that a round exists in which p progresses, as do a set of other correct participants who have sufficient availability for p to continue to progress, and all of them send messages of identical value.

By requirement 4, this is sufficient for p to decide that value, provided p receives all of those messages first. This will occur with some non-zero probability. Therefore, in any pair of consecutive rounds, there is a non-zero probability a guru will decide. Therefore, each guru, with probability 1, eventually decides.

References

[1] I. Abraham, D. Dolev, R. Gonen, and J. Halpern. Distributed computing meets game theory: Robust mechanisms for rational secret sharing and multiparty computation. In *Proceedings of the Twenty-fifth Annual ACM Symposium on Principles of Distributed Computing*, PODC '06, pages 53–62, New York, NY, USA, 2006. ACM.

[2] Y. Afek, H. Attiya, A. Fekete, M. Fischer, N. Lynch, Y. Mansour, D.-W. Wang, and L. Zuck. Reliable communication over unreliable channels. *J. ACM*, 41(6):1267–1297, Nov. 1994. ISSN 0004-5411. . URL http://doi.acm.org/10.1145/195613.195651.

[3] A. S. Aiyer, L. Alvisi, A. Clement, M. Dahlin, J.-P. Martin, and C. Porth. BAR fault tolerance for cooperative services. In *Proceedings of the Twentieth ACM Symposium on Operating Systems Principles*, SOSP '05, pages 45–58, New York, NY, USA, 2005. ACM. ISBN 1-59593-079-5. . URL http://doi.acm.org/10.1145/1095810.1095816.

[4] K. J. Biba. Integrity considerations for secure computer systems. Technical Report ESD-TR-76-372, USAF Electronic Systems Division, Bedford, MA, Apr. 1977. (Also available through National Technical Information Service, Springfield Va., NTIS AD-A039324.).

[5] R. S. Boyer and J. S. Moore. A mechanical proof of the unsolvability of the halting problem. *J. ACM*, 31(3):441–458, June 1984. ISSN 0004-5411. . URL http://doi.acm.org/10.1145/828.1882.

[6] G. Bracha and S. Toueg. Asynchronous consensus and broadcast protocols. *Journal of the ACM*, 32(4):824–240, 1985.

[7] F. Brasileiro, F. Greve, A. Mostefaoui, and M. Raynal. Consensus in one communication step. In V. Malyshkin, editor, *Parallel Computing Technologies*, volume 2127 of *Lecture Notes in Computer Science*, pages 42–50. Springer Berlin Heidelberg, 2001. ISBN 978-3-540-42522-9. . URL http://dx.doi.org/10.1007/3-540-44743-1_4.

[8] M. Castro and B. Liskov. Practical Byzantine fault tolerance. In *Proc. 3rd Symposium on Operating Systems Design and Implementation*, New Orleans, LA, Feb. 1999.

[9] S. Chong. *Expressive and Enforceable Information Security Policies*. PhD thesis, Cornell University, Aug. 2008.

[10] S. Chong and A. C. Myers. Decentralized robustness. In *Proc. 19th IEEE Computer Security Foundations Workshop*, pages 242–253, July 2006.

[11] S. N. Foley. A taxonomy for information flow policies and models. In *Proc. IEEE Symp. on Security and Privacy*, pages 98–108, 1991.

[12] J. Garay and K. Perry. A continuum of failure models for distributed computing. In A. Segall and S. Zaks, editors, *Distributed Algorithms*, volume 647 of *Lecture Notes in Computer Science*, pages 153–165. Springer Berlin Heidelberg, 1992. ISBN 978-3-540-56188-0. . URL http://dx.doi.org/10.1007/3-540-56188-9_11.

[13] C. Ho, D. Dolev, and R. Renesse. Making distributed applications robust. In E. Tovar, P. Tsigas, and H. Fouchal, editors, *Principles of Distributed Systems*, volume 4878 of *Lecture Notes in Computer Science*, pages 232–246. Springer Berlin Heidelberg, 2007. ISBN 978-3-540-77095-4. . URL http://dx.doi.org/10.1007/978-3-540-77096-1_17.

[14] C. Ho, R. Van Renesse, M. Bickford, and D. Dolev. Nysiad: Practical protocol transformation to tolerate Byzantine failures. In *Proc. 5th USENIX Symp. on Networked Systems Design and Implementation (NSDI)*, volume 8, pages 175–188, 2008.

[15] A. Jaffe, T. Moscibroda, and S. Sen. On the price of equivocation in Byzantine agreement. In *Proceedings of the 2012 ACM Symposium on Principles of Distributed*

Computing, PODC '12, pages 309–318, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1450-3. . URL http://doi.acm.org/10.1145/2332432.2332491.

[16] F. Junqueira and K. Marzullo. Designing algorithms for dependent process failures. In *Proceedings of the Workshop on Future Directions in Distributed Computing*, pages 24–28, 2003.

[17] F. Junqueira and K. Marzullo. Replication predicates for dependent-failure algorithms. In J. Cunha and P. Medeiros, editors, *Euro-Par 2005 Parallel Processing*, volume 3648 of *Lecture Notes in Computer Science*, pages 617–632. Springer Berlin Heidelberg, 2005. ISBN 978-3-540-28700-1. . URL http://dx.doi.org/10.1007/11549468_69.

[18] L. Lamport. The Part-time Parliament. *ACM Trans. Comput. Syst.*, 16(2):133–169, May 1998. ISSN 0734-2071. . URL http://doi.acm.org/10.1145/279227.279229.

[19] L. Lamport. Fast Paxos. *Distributed Computing*, 19(2):79–103, October 2006. URL http://research.microsoft.com/apps/pubs/default.aspx?id=64624.

[20] L. Lamport, R. Shostak, and M. Pease. The Byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, July 1982. ISSN 0164-0925. . URL http://doi.acm.org/10.1145/357172.357176.

[21] B. Lampson, M. Abadi, M. Burrows, and E. Wobber. Authentication in distributed systems: Theory and practice. In *Proc. 13th ACM Symp. on Operating System Principles (SOSP)*, pages 165–182, Oct. 1991. *Operating System Review*, 253(5).

[22] D. Malkhi and M. Reiter. Byzantine quorum systems. In *Proc. 29th ACM Symposium on Theory of Computing*, pages 569–578, El Paso, Texas, May 1997.

[23] F. Meyer and D. Pradhan. Consensus with dual failure modes. *Parallel and Distributed Systems, IEEE Transactions on*, 2(2):214–222, Apr 1991. ISSN 1045-9219. .

[24] A. C. Myers and B. Liskov. Protecting privacy using the decentralized label model. *ACM Transactions on Software Engineering and Methodology*, 9(4):410–442, Oct. 2000. URL http://www.cs.cornell.edu/andru/papers/iflow-tosem.pdf.

[25] A. C. Myers, L. Zheng, S. Zdancewic, S. Chong, and N. Nystrom. Jif 3.0: Java information flow. Software release, http://www.cs.cornell.edu/jif, July 2006.

[26] J. Palsberg and P. Ørbæk. Trust in the λ-calculus. In *Proc. 2nd International Symposium on Static Analysis*, number 983 in Lecture Notes in Computer Science, pages 314–329. Springer, Sept. 1995.

[27] J. Postel. DoD standard transmission control protocol. DARPA-Internet RFC-793, Sept. 1981.

[28] A. Sabelfeld and A. C. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 21(1):5–19, Jan. 2003. URL http://www.cs.cornell.edu/andru/papers/jsac/sm-jsac03.pdf.

[29] R. D. Schlichting and F. B. Schneider. Fail-stop processors: An approach to designing fault-tolerant computing systems. *ACM Transactions on Computer Systems*, 1(3):222–238, 1983.

[30] F. B. Schneider. Implementing fault-tolerant services using the state machine approach: a tutorial. *ACM Computing Surveys*, 22(4):299–319, Dec. 1990.

[31] H.-S. Siu, Y.-H. Chin, and W.-P. Yang. Byzantine agreement in the presence of mixed faults on processors and links. *Parallel and Distributed Systems, IEEE Transactions on*, 9(4):335–345, Apr 1998. ISSN 1045-9219. .

[32] Y. J. Song and R. Renesse. Bosco: One-step Byzantine asynchronous consensus. In *Proc. 22nd Int'l Symp. on Distributed Computing*, DISC '08, pages 438–450,

Berlin, Heidelberg, 2008. Springer-Verlag. ISBN 978-3-540-87778-3. . URL http://dx.doi.org/10.1007/978-3-540-87779-0_30.

[33] P. Steiner. On the Internet, nobody knows you're a dog. *The New Yorker*, 69(20):61, 1993. URL http://www.unc.edu/depts/jomc/academics/dri/idog.html.

[34] D. Walker, L. Mackey, J. Ligatti, G. Reis, and D. August. Static typing for a faulty lambda calculus. In *ACM SIGPLAN International Conference on Functional Programming*, Sept. 2006.

[35] S. Zdancewic, L. Zheng, N. Nystrom, and A. C. Myers. Secure program partitioning. *ACM Transactions on Computer Systems*, 20(3):283–328, Aug. 2002. URL http://www.cs.cornell.edu/andru/papers/sosp01/spp-tr.pdf.

[36] D. Zhang and A. C. Myers. Toward general diagnosis of static errors. In *Proc. 41st ACM Symposium on Principles of Programming Languages (POPL)*, pages 569–581, Jan. 2014. URL http://www.cs.cornell.edu/andru/papers/diagnostic.

[37] L. Zheng. *Making distributed computation secure by construction*. PhD thesis, Cornell University, Ithaca, New York, USA, Jan. 2007.

[38] L. Zheng and A. C. Myers. End-to-end availability policies and noninterference. In *Proc. 18th IEEE Computer Security Foundations Workshop*, pages 272–286, June 2005.

[39] L. Zheng and A. C. Myers. A language-based approach to secure quorum replication. In *Proceedings of the Ninth Workshop on Programming Languages and Analysis for Security*, PLAS'14, pages 27:27–27:39, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2862-3. . URL http://doi.acm.org/10.1145/2637113.2637117.

[40] L. Zheng, S. Chong, A. C. Myers, and S. Zdancewic. Using replication and partitioning to build secure distributed systems. In *Proc. IEEE Symp. on Security and Privacy*, pages 236–250, May 2003. URL http://www.cs.cornell.edu/andru/papers/sp03.pdf.